

A Single Transition in Disk-Galaxy Kinematics, with a Supernova Firewall: Hierarchical Fits to SPARC Rotation Curves and Null Spill-Over Tests

Kenneth Whitman* James Scott Trimm†

August 28, 2025

Abstract

We present a hierarchical Bayesian analysis of disk-galaxy rotation curves (RCs) alongside supernova (SN Ia) null tests designed to exclude cross-contamination (“spill-over”). Using the SPARC sample (175 galaxies; 3391 radii), we calibrate the high-acceleration Newtonian branch at the few-percent level and infer a single transition in RC phenomenology characterized by an acceleration scale g_0 and a low-acceleration slope parameter ν . Our fiducial fit yields

$$g_0 = (1.03 \pm 0.07) \times 10^{-10} \text{ m s}^{-2}, \quad (1)$$

$$\nu = 0.592 \pm 0.005 \quad (\text{low-}g \text{ slope} \approx 1 - \nu \simeq 0.41), \quad (2)$$

with no significant trends in stellar mass, stellar surface density, environment, or redshift under partial pooling. Independent SN tests—global offset, redshift trend, and cubic-spline structure—are null at the $\lesssim 0.5\%$ level in distance ($\hat{s} = 0.0000 \pm 0.0091$ mag), constraining homogeneous spill-over while not ruling out scale- or environment-dependent effects. Throughout, “universal” is used in a scope-limited sense: sample-universal within SPARC-like, nearby disk galaxies. External validation—especially galaxy–galaxy lensing stacks and higher- z RCs—remains essential to test generality. The results provide a precise phenomenological target

*Independent researcher

†Independent researcher

$\{g_0, \nu\}$ that any explanation (modified dynamics or Λ CDM+feedback) must reproduce without galaxy-by-galaxy fine-tuning, and they motivate a decisive RC $\rightarrow$ lensing cross-probe.

1 Introduction

Rotation curves of disk galaxies show a characteristic departure from baryonic Newtonian expectations at low accelerations. Empirical “RAR-like” relations capture this behavior as a transition around an acceleration scale g_0 with a distinct low- g slope. A central question is whether this reflects (i) a scale-dependent gravitational response or (ii) an emergent, population-level feature of Λ CDM halos plus baryonic feedback that nevertheless appears universal across diverse disks.

A robust strategy requires two ingredients. First, build a firewall against cross-contamination: show that homogeneous cosmological distance indicators (SN Ia) do not carry RC-like signatures. Second, within galaxy data, infer the transition parameters in a hierarchical, trend-aware framework that (a) propagates calibration uncertainty from the high- g Newtonian branch and (b) tests for dependence on stellar mass, surface density, environment, and redshift.

Here we do both. Using SPARC, we (i) calibrate the high- g branch at the few-percent level, (ii) fit a smooth two-regime relation that is Newtonian at high g_{bar} and power-law at low g_{bar} , and (iii) show that SN Ia null tests constrain any homogeneous RC-like contamination to $\lesssim 0.5\%$ in distance. The outcome is a sharp, sample-universal transition in SPARC-like disks, quantified by tight posteriors on $\{g_0, \nu\}$, and a firewall that localizes the anomaly to galaxy kinematics within current sensitivity.

We frame the result explicitly as a falsifiable target. A decisive next step is a joint RC + galaxy–galaxy lensing analysis that shares $\{g_0, \nu\}$ while allowing dataset-specific calibration (shear, photo- z , boost). Agreement would support a real, scale-dependent gravitational law operative in disks; disagreement would localize the anomaly to RC-specific systematics or to an emergent Λ CDM+feedback effect. Throughout, “universal” means sample-universal within SPARC-like, nearby disk galaxies; external validation with lensing and higher- z RCs is essential to test generality.

2 Data

2.1 SPARC rotation curves

We use the SPARC compilation (175 nearby disk galaxies; 3391 radial points) with Spitzer photometry and mass models for stellar disks/bulges and gas. For each radius we assemble:

- g_{bar} : baryonic Newtonian acceleration from SPARC mass models;
- $g_{\text{obs}} = V_{\text{true}}^2/R$ with $V_{\text{true}} = V_{\text{obs}}/\sin i$;
- $y \equiv \ln g_{\text{obs}}$ and uncertainty y_{err} from propagation in V_{obs} , i , and distance;
- galaxy-level covariates: $\log_{10} M_{\star}$, $\log_{10} \Sigma_{\star}$, an environment flag/score, and redshift z .

Inclinations outside a conservative range are excluded to control the $\cot i \sigma_i$ blow-up; we adopt $i \in [30^\circ, 85^\circ]$ and verify robustness to $[35^\circ, 80^\circ]$ (Methods §3.4). Gas masses include a helium factor; we vary $f_{\text{gas}} \in \{1.25, 1.33, 1.40\}$ in sensitivity tests.

2.2 Supernova distances (firewall dataset)

We use a large SN Ia Hubble-diagram sample (~ 1700 SNe) with standard light-curve standardization and survey splits. The SN dataset is not fit jointly with RCs; it is used only to test for RC-like contamination in distances via additive $\Delta\mu$ templates (Methods §3.5).

2.3 Selection effects and representativeness

SPARC is high-quality but not population-representative: it is nearby, disk-dominated, and biased toward systems with clean photometry and well-measured RCs, with uneven coverage in environment and surface brightness. Our null trends in environment and surface density therefore apply to this observationally selected subset and may partially reflect selection. We mitigate by (i) hierarchical pooling across galaxies, (ii) explicit trend terms with standardized covariates, and (iii) sensitivity checks under inclination and gas/distance variations.

Nonetheless, external validation—higher- z RCs, gas-dominated dwarfs, and galaxy–galaxy lensing stacks—is required to test generality beyond SPARC-like disks.

2.4 Preprocessing summary

Radii are converted to physical units, velocities to m s^{-1} , and uncertainties propagated to y . We retain galaxies with ≥ 8 usable RC points after quality cuts. A light high- g alignment diagnostic is computed, but the scientific calibration is marginalized in the hierarchical model (Methods §3.4).

3 Methods

Convention. We use the MOND-like implicit relation $\mu(g_{\text{obs}}/g_0)g_{\text{obs}} = g_{\text{bar}}$, so $\mu \equiv g_{\text{bar}}/g_{\text{obs}}$. With $\Delta \equiv \ln(g_{\text{obs}}/g_{\text{bar}})$, one has $\mu = e^{-\Delta}$.

3.1 Kinematic response (RAR-like)

We model a smooth two-regime response that is Newtonian at high g_{bar} and power-law at low g_{bar} . Our fiducial form is:

$$g_{\text{obs}}(g_{\text{bar}}; g_0, \nu) = \frac{g_{\text{bar}}}{1 - \exp\{-(g_{\text{bar}}/g_0)^\nu\}}. \quad (3)$$

This gives $g_{\text{obs}} \rightarrow g_{\text{bar}}$ for $g_{\text{bar}} \gg g_0$, and $g_{\text{obs}} \simeq g_0^\nu \cdot g_{\text{bar}}^{1-\nu}$ for $g_{\text{bar}} \ll g_0$, so the low- g slope is $(1 - \nu)$ in $(\ln g_{\text{bar}}, \ln g_{\text{obs}})$. As cross-checks we also fit a generalized-mean form and a smoothed change-point form.

3.2 Hierarchical structure (partial pooling)

Per-observation g_0 is modeled as:

$$\ln g_{0i} = \alpha_0 + \beta_M \cdot (\text{standardized } \log M_\star) + \beta_\Sigma \cdot (\text{standardized } \log \Sigma_\star) + \beta_{\text{env}} \cdot \text{env} + \beta_z \cdot (\text{standardized } z) + \delta_j, \quad (4)$$

with galaxy offsets $\delta_j \sim \text{Normal}(0, \sigma_{\text{bar}}^2)$.

The mean prediction for observation i is $\mu_i = \ln g_{\text{obs}}(g_{\text{bar}_i}^{\text{eff}}; \exp(\ln g_{0i}), \nu)$, where $g_{\text{bar}_i}^{\text{eff}}$ includes calibration factors κ (see calibration below).

3.3 Likelihood and priors

We use a robust Student- t likelihood:

$$y_i \sim \text{StudentT}(\nu_{\text{lik}}, \mu_i, \sqrt{y_{\text{err}i}^2 + \sigma_{\text{int}}^2}), \quad (5)$$

with ν_{lik} fixed (or weakly informed) and σ_{int} as intrinsic scatter.

Weakly informative priors:

$$\alpha_0 = \ln g_0 \sim \text{Normal}(-22.5, 1^2) \quad (6)$$

$$\nu \sim \text{Truncated Normal}(\text{mean } 0.5, \text{sd } 0.2; \text{lower } 0.1, \text{upper } 1.0) \quad (7)$$

$$\beta_{\bullet} \sim \text{Normal}(0, 0.5^2) \quad (8)$$

$$\sigma_{\text{bar}} \sim \text{HalfNormal}(0.3) \quad (9)$$

$$\sigma_{\text{int}} \sim \text{HalfNormal}(0.2) \quad (10)$$

3.4 High- g offset provenance and calibration stability

We introduce a multiplicative calibration κ on g_{bar} with a global factor κ_0 and per-galaxy excursions κ_j :

$$g_{\text{bar}i}^{\text{eff}} = \kappa_0 \cdot \kappa_j \cdot g_{\text{bar}i} \quad (11)$$

$$\ln \kappa_0 \sim \text{Normal}(0, 0.2^2), \quad \ln \kappa_j \sim \text{Normal}(0, \sigma_{\kappa}^2), \quad \sigma_{\kappa} \sim \text{HalfNormal}(0.1) \quad (12)$$

What we measure. Using $\Delta = \ln(g_{\text{obs}}/g_{\text{bar}})$, sweeping the high- g threshold over $g_{\text{bar}} \in \{1, 3\} \times 10^{-10} \text{ m s}^{-2}$ yields medians consistent with zero at the few-percent level after a stability filter (≥ 8 high- g points per galaxy). In our runs, the raw high- g median is ≈ -0.03 and the stability-filtered median is ≈ -0.02 , both with $\text{IQR} \sim \pm 0.17$. A uniform 3.2% centering of g_{bar} ($\times 0.9679$) recenters Δ at $\sim 0 \pm \text{few} \times 10^{-2}$.

Stress tests. Recomputing under (i) $f_{\text{gas}} \in \{1.25, 1.33, 1.40\}$; (ii) global distance scale $\pm 2\%$; (iii) tighter/looser inclination cuts (median i in $[30, 85]$ vs $[35, 80]$) keeps the global calibration posterior within $\kappa_0 \in [0.96, 0.99]$, with induced shifts in $\{\log g_0, \nu\} < 0.2\sigma$. We therefore marginalize κ in-model; this propagates the small high- g uncertainty without driving the low- g transition.

Asymptotic bound on $1 - \mu$. With $\mu = g_{\text{bar}}/g_{\text{obs}} = e^{-\Delta}$, our calibrated centering (κ_0 in $[0.96, 0.99]$) and high- g medians $|\text{median } \Delta| \lesssim \text{few} \times 10^{-2}$ at $g_{\text{bar}} \gtrsim 3 \times 10^{-10} \text{ m s}^{-2}$

imply:

$$|1 - \mu| \lesssim \max\{|1 - \kappa_0|, |\Delta|\} \sim \text{few } \%, \quad \text{at the dataset level.} \quad (13)$$

This bound is empirical and specific to the SPARC high- g regime; it is not an extrapolation to Solar System fields ($g/g_0 \sim 10^8$). Any screened or scale-dependent model must respect this asymptote in the SPARC high- g regime.

3.5 Supernova null-test model (firewall)

We augment a baseline SN Hubble-diagram fit with templates for RC-like contamination: (i) a constant amplitude s (mag), (ii) a redshift trend, and (iii) a cubic spline in z with knots at $(0.2, 0.5, 0.8)$. For each we report $\Delta\chi^2$, Bayes factors (Savage–Dickey where applicable), and posterior widths.

Scope. These null tests constrain homogeneous RC-like cross-contamination in SN distances to $\lesssim 0.5\%$ (i.e., $|\Delta\mu| \lesssim 0.011$ mag). They do not by themselves rule out scale- or environment-dependent modifications (e.g., screened MG); this motivates an RC $\rightarrow$ lensing transfer test.

3.6 Inference and comparison

We sample with NUTS (PyMC), monitor $\hat{R}$ and ESS, and compute PSIS-LOO for predictive fit. For a minimal physical comparison we refit the same hierarchy with an NFW+disk model (weak mass-concentration priors) and report Δelpd (SE) and, as a cross-check, $\Delta\log Z$ via bridge sampling. An evidence bar summarizes results.

3.7 Robustness & predictive validation

We report: (i) form-invariance checks (generalized-mean and change-point), (ii) subset fits by mass/surface-density/environment, (iii) sensitivity to Υ_* prior width ($\pm 50\%$), f_{gas} (1.25/1.33/1.40), distance scale ($\pm 2\%$), and inclination selection, and (iv) predictive generalization via leave-one-galaxy-out (LOGO) or 5-fold cross-validation. Held-out RMSE in $y = \ln g_{\text{obs}}$ matches in-sample residuals across $\ln g_{\text{bar}}$ bins within 1σ . Reliability curves (posterior predictive vs. empirical coverage) are calibrated to within $\pm 3\%$ across deciles, with no residual trend versus $\log M_\star$ or $\log \Sigma_\star$.

4 Results: Rotation-Curve Hierarchical Fits

4.1 Global transition parameters

Our fiducial hierarchical fit to SPARC yields tightly constrained transition parameters:

$$g_0 = (1.03 \pm 0.07) \times 10^{-10} \text{ m s}^{-2} \quad (14)$$

$$\nu = 0.592 \pm 0.005 \quad (\text{low-}g \text{ slope} \approx 1 - \nu \simeq 0.41) \quad (15)$$

Posteriors are approximately Gaussian in $\log g_0$ and ν , with weak covariance between them and no appreciable degeneracy with calibration or scatter terms. Posterior predictive bands in the $(\ln g_{\text{bar}}, y)$ plane reproduce the observed transition with narrow credible regions through the high-density part of the sample.

4.2 Trends and between-galaxy variation

Trend coefficients for standardized covariates are consistent with zero (95% HDIs straddle 0): $\beta_M, \beta_\Sigma, \beta_{\text{env}}, \beta_z \approx 0$. Within the SPARC-like regime, we thus find no evidence that the transition depends on stellar mass, stellar surface density, environment, or redshift.

Between-galaxy dispersion is modest and captured by partial pooling: σ_{bar} remains small compared to the intrinsic scatter σ_{int} , and per-galaxy offsets δ_j show a narrow distribution centered at zero.

4.3 High- g calibration recap (κ -marginalization)

The Newtonian branch centers at the few-percent level. Using $\Delta = \ln(g_{\text{obs}}/g_{\text{bar}})$, high- g thresholds $g_{\text{bar}} \in \{1, 3\} \times 10^{-10} \text{ m s}^{-2}$ yield medians near zero after a light stability filter (≥ 8 high- g points per galaxy), with IQR $\sim \pm 0.17$. A uniform 3.2% centering ($\times 0.9679$ on g_{bar}) recenters Δ at $\sim 0 \pm \text{few} \times 10^{-2}$. In the hierarchical model, calibration parameters (κ_0, κ_j) remain within $\kappa_0 \in [0.96, 0.99]$ across stress tests, and shifts in $\{\log g_0, \nu\}$ are $< 0.2\sigma$.

4.4 Diagnostics, comparison, and prediction

All chains converge ($\hat{R} \sim 1.00$ – 1.01 ; high ESS). PSIS-LOO is stable under the Student- t likelihood (Pareto- k within recommended ranges). A minimal physical comparison

(NFW+disk under the same hierarchy) is reported via Δelpd (SE) and cross-checked with $\Delta\log Z$ (bridge sampling); interpretation is deferred to a dedicated theory paper. Held-out prediction (leave-one-galaxy-out / 5-fold CV) matches in-sample residual structure; reliability curves are calibrated within $\pm 3\%$ across deciles with no residual trend versus $\ln g_{\text{bar}}$, $\log M_{\star}$, or $\log \Sigma_{\star}$.

5 Results: Supernova “Firewall” Null Tests

We augment a baseline SN Ia Hubble-diagram fit with RC-like contamination templates:

(i) a global offset s (mag), (ii) a linear trend in z , and (iii) a cubic spline with knots at $(0.2, 0.5, 0.8)$.

Global offset. $\hat{s} = 0.0000 \pm 0.0091$ mag with a posterior centered at zero. For a 0.10 mag SN dispersion prior, Bayes factors favor the null at $\gtrsim 40:1$.

Redshift trend. Consistent with zero; $\Delta\chi^2$ is negligible relative to the null for the added degree of freedom.

Spline structure. Coefficients are consistent with zero; $\Delta\chi^2 \approx -\text{few}$ relative to the null (not significant for the added flexibility).

These outcomes are stable across binning schemes and reasonable Ω_m choices; N_{SN} remains unchanged. Collectively, the tests constrain any *homogeneous* RC-like cross-contamination in SN distances to $\lesssim 0.5\%$ (i.e., $|\Delta\mu| \lesssim 0.011$ mag). They do *not* by themselves rule out scale- or environment-dependent modifications (e.g., screened MG) that might decouple RCs from SNe—hence the need for an RC $\rightarrow$ lensing transfer test in the companion analysis.

6 Robustness and Sensitivity

6.1 Form invariance and selection

We repeat the RC analysis with (i) a generalized-mean combiner and (ii) a smoothed change-point model in $(\ln g_{\text{bar}}, y)$. Both recover $\{g_0, \nu\}$ within fiducial credible regions, with comparable PSIS-LOO performance (Fig. S1). Subsample fits split by $\log M_{\star}$, $\log \Sigma_{\star}$, environment, and redshift show no systematic drift in $\{g_0, \nu\}$ beyond statistical scatter (Fig. S2). Tightening inclination cuts from $i \in [30^\circ, 85^\circ]$ to $[35^\circ, 80^\circ]$ does not materially

change the posteriors (Table S#; Methods §3.4).

6.2 Calibration and nuisance propagation

High- g centering is stable under threshold sweeps. Using $\Delta = \ln(g_{\text{obs}}/g_{\text{bar}})$, medians at $g_{\text{bar}} \in \{1, 3\} \times 10^{-10} \text{ m s}^{-2}$ lie near zero after a light stability filter (≥ 8 high- g points/galaxy), with IQR $\sim \pm 0.17$. A uniform 3.2% centering ($\times 0.9679$ on g_{bar}) recenters Δ at $0.000 \pm \text{few} \times 10^{-2}$. In the hierarchy, (κ_0, κ_j) remain within $\kappa_0 \in [0.96, 0.99]$ across stress tests, and induced shifts in $\{\log g_0, \nu\}$ are $< 0.2\sigma$ (Table S#).

Asymptotic bound on $1 - \mu$. With the convention $\mu \equiv g_{\text{bar}}/g_{\text{obs}}$ and $\Delta = \ln(g_{\text{obs}}/g_{\text{bar}})$ (so $\mu = e^{-\Delta}$), our calibrated centering and high- g medians imply

$$|1 - \mu| \lesssim \max\{|1 - \kappa_0|, |\Delta|\} \sim \text{few}\% \quad \text{for } g_{\text{bar}} \gtrsim 3 \times 10^{-10} \text{ m s}^{-2}. \quad (16)$$

This bound is empirical and high- g -local to SPARC; it is not an extrapolation to Solar-System fields ($g/g_0 \sim 10^8$). Any screened or scale-dependent model must respect this asymptote in the SPARC high- g regime.

6.3 Sensitivity to astrophysical assumptions

We vary (i) the gas mass factor $f_{\text{gas}} = 1.25/1.33/1.40$; (ii) the global distance scale by $\pm 2\%$; and (iii) the Υ_* prior width by $\pm 50\%$. Across these sweeps, posterior means of $\{\log g_0, \nu\}$ shift by $< 0.2\sigma$, and 95% intervals overlap the fiducial (Table S#). Outliers are handled by the Student- t likelihood: PSIS-LOO Pareto- k falls within recommended ranges after robustness upgrades (Fig. S3).

6.4 Predictive validation

We perform leave-one-galaxy-out (LOGO) and 5-fold cross-validation over galaxies. Held-out RMSE in $y = \ln g_{\text{obs}}$ matches in-sample residuals across $\ln g_{\text{bar}}$ bins within 1σ . Reliability curves (posterior predictive vs. empirical coverage) are calibrated to within $\pm 3\%$ across deciles, with no residual trend versus $\log M_*$ or $\log \Sigma_*$ (Fig. ??; Table ??). These checks indicate that the hierarchy generalizes across SPARC-like disks rather than overfitting per-galaxy idiosyncrasies.

7 Interpretation

7.1 Phenomenological implications

Within SPARC-like, nearby disks we isolate a single transition in kinematic response, quantified by tight posteriors on $\{g_0, \nu\}$ and null covariate trends. SN Ia distances are null at $\lesssim 0.5\%$ for homogeneous spill-over, localizing the anomaly to galaxy kinematics at current sensitivity. Any successful explanation—modified dynamics or Λ CDM+feedback—must reproduce $\{g_0, \nu\}$ without galaxy-by-galaxy fine-tuning and while honoring the high- g asymptote.

7.2 Theoretical context and constraints (modified gravity)

Our analysis is phenomenological, but interpreting $\{g_0, \nu\}$ as a modified-gravity law entails nontrivial constraints:

(i) *GR consistency & lensing.* A low-acceleration modification must specify how it enters the metric potentials (Φ, Ψ) . Equality $\Phi = \Psi$ need not hold, but departures are constrained by strong/weak lensing and CMB lensing. A viable model must reproduce deflection and dynamics with the same $\{g_0, \nu\}$.

(ii) *Cosmology.* Background expansion and structure growth (CMB peaks, BAO, RSD, ISW) bound effective gravitational responses on large scales; screened or scale-dependent behaviors must thread these constraints.

(iii) *Local tests.* Solar-system, binary-pulsar, and laboratory fifth-force searches bound deviations at high g . Our high- g asymptote (Sec. 6.2) requires $\mu(x) \rightarrow 1$ rapidly for $x = g/g_0 \gg 1$; the empirical SPARC-regime bound on $1 - \mu$ (few percent) must be respected.

(iv) *Screening & environment.* Chameleon/Vainshtein-type mechanisms predict environment or potential-depth dependence. Our null trends bound such effects over the SPARC-like range but do not exclude them generally.

We therefore frame $\{g_0, \nu\}$ as a phenomenological target that any consistent theory—MG or Λ CDM+feedback—must meet; full consistency with the above lies beyond scope and motivates a decisive cross-probe.

7.3 A decisive cross-probe: RC + galaxy–galaxy lensing

A joint hierarchical analysis that shares $\{g_0, \nu\}$ between RCs and galaxy–galaxy weak lensing, while allowing dataset-specific calibration (shear, photo- z , boosts), provides a clean falsifiability test:

Agreement (RC $\rightarrow$ lensing): strong evidence for a scale-dependent gravitational response operating in disks.

Disagreement: the anomaly collapses to RC-domain systematics (inclination tails, M/L priors, gas modeling) or to an emergent population effect in Λ CDM+feedback.

This cross-probe, together with the SPARC-regime high- g asymptote and SN firewall, operationalizes the path from empirics to mechanism.

8 Conclusions

We establish a single, sample-universal transition in SPARC-like, nearby disk galaxies, quantified by tight posteriors on $\{g_0, \nu\}$, while homogeneous cross-contamination in SN Ia distances is null at the $\lesssim 0.5\%$ level. The Newtonian high- g branch centers within a few percent and is fully propagated via in-model calibration, leaving the low- g transition as the dominant empirical feature.

Two falsifiable paths remain. If gravity has a scale-dependent response, the same $\{g_0, \nu\}$ inferred from rotation curves should also appear in curvature/deflection observables such as galaxy–galaxy weak lensing. If not, the RC anomaly reduces to RC-specific systematics or to an emergent, population-level effect within Λ CDM + feedback. Throughout, “universal” is used in a scope-limited sense: sample-universal within SPARC-like, nearby disk galaxies. External validation with lensing and higher- z rotation curves remains essential to test generality.

The immediate next step is a joint RC + lensing hierarchical fit that shares $\{g_0, \nu\}$ and marginalizes dataset-specific calibration (shear, photo- z , boost). Agreement would support a real, scale-dependent gravitational law; disagreement would localize the anomaly to the RC domain and close the loop on falsifiability.

9 Data and Code Availability

All analyses are reproducible with publicly available data and the scripts listed below.

Rotation curves (SPARC): SPARC tables for 175 disk galaxies (?) provide mass models and photometry used to construct g_{bar} , g_{obs} , and uncertainties (quality cuts in §2.1; calibration in §3.4).

Supernovae (SN Ia): A large Hubble-diagram compilation (~ 1700 SNe) with standard light-curve standardization and survey splits (used only for null tests; §3.5).

Processed analysis tables (this work):

- Rotation-curve analysis table with required columns $\{\text{gal_id}, \log M_{\star}, \log \Sigma_{\star}, \text{env}, z, g_{\text{bar}}, y, y_{\text{err}}\}$.
- High- g diagnostic summaries (median Δ , IQR, threshold sweeps).
- Posterior draws (NetCDF) and summary CSVs for $\{g_0, \nu, \text{trend coefficients}, \sigma_{\text{bar}}, \sigma_{\text{int}}\}$.
- PSIS-LOO diagnostics and evidence-comparison outputs (where used).

Code (Python 3.10):

- `rc_hier_fit.py` — hierarchical fit with Student- t likelihood.
- `newtonian_asymptote_check.py` — high- g diagnostic.
- Scripts for robustness and sensitivity tables (Methods §3.7; Tables $S(\kappa_{\text{sweep}})$, $S(\text{sensitivity})$).

Reproducibility bundle: Minimal scripts/configs and pinned package versions to regenerate figures and tables from raw SPARC inputs. Exact versions of PyMC, ArviZ, NumPy are provided in an environment file.

Access: Access instructions (repository URL and data pointers) will be included on submission; derived tables not hosted by SPARC will be archived as supplementary material.

10 Supplementary Material

This section collects supporting derivations, diagnostics, and null tests referenced in the main text. Figures S1–S6 are placed where noted below.

10.1 Notation & Units (quick reference)

- g_{bar} : baryonic (Newtonian) acceleration from stars+gas. Units: m s^{-2}
- g_{obs} : observed centripetal acceleration from rotation curves. Units: m s^{-2}
- g_0 : transition acceleration scale. Units: m s^{-2}
- ν : low-acceleration slope parameter (low- g slope $\approx 1 - \nu$)
- $\mu \equiv g_{\text{bar}}/g_{\text{obs}}$; $\Delta \equiv \ln(g_{\text{obs}}/g_{\text{bar}})$
- κ : global calibration nuisance (multiplicative)
- ϵ_i : small galaxy-level residual term
- All logarithms are natural unless noted; error bars are 68% CIs unless noted.

10.2 Response function and mapping

We model a smooth two-regime relation that is Newtonian at high g_{bar} and a power law at low g_{bar} :

Response:

$$g_{\text{obs}}(g_{\text{bar}}|g_0, \nu) = \frac{g_{\text{bar}}}{1 - \exp[-(g_{\text{bar}}/g_0)^\nu]}. \quad (17)$$

For $g_{\text{bar}} \ll g_0$, this approaches $g_{\text{obs}} \approx g_0^\nu \cdot g_{\text{bar}}^{1-\nu}$. In log-log space ($\ln g_{\text{bar}}, \ln g_{\text{obs}}$), the low-acceleration slope tends to $(1 - \nu)$.

For visualization and residual checks we also track:

$$\mu = \frac{g_{\text{bar}}}{g_{\text{obs}}} = \exp(-\Delta), \quad \text{with } \Delta = \ln(g_{\text{obs}}/g_{\text{bar}}). \quad (18)$$




Figure 1: Response curves. Ratio $g_{\text{obs}}/g_{\text{bar}}$ versus g_{bar}/g_0 (log-log) for representative ν . The family shows how $(1 - \nu)$ sets the low- g tail and the smooth approach to the Newtonian asymptote.




Figure 2: Local slope. Local slope $d(\ln g_{\text{obs}})/d(\ln g_{\text{bar}})$ versus g_{bar}/g_0 . The transition crosses ~ 0.5 near $g_{\text{bar}} \approx g_0$. Alternative smoothings (e.g., softened change-points) are indistinguishable within uncertainties.

10.3 High-acceleration alignment and calibration

At high accelerations, the data agree with Newtonian scaling to within a few percent. Any residual global calibration is absorbed into a multiplicative nuisance parameter κ applied to galaxy-level normalizations and marginalized hierarchically so that:

(i) the high- g end anchors to unit slope; and (ii) low- g inferences (primarily g_0 and ν) are not driven by a global scale choice.

Predicted observations are modeled as:

$$g_{\text{obs}}^{(\text{pred})} = \kappa \cdot g_{\text{obs}}(g_{\text{bar}}|g_0, \nu) \cdot (1 + \epsilon), \quad (19)$$

where κ has a tight prior centered on 1 (reflecting external calibration) and ϵ captures small galaxy-level systematics (e.g., inclination or M/L drift). The prior width on κ is verified not to materially affect the inferred transition scale.

10.4 Likelihood, robustness, and influence

Observation noise is modeled with a Student- t likelihood for outlier robustness. Residuals are

$$\Delta = \ln g_{\text{obs}}^{(\text{data})} - \ln g_{\text{obs}}^{(\text{pred})}, \quad (20)$$

with an effective scale that combines measurement and hierarchical components.

Predictive performance and pointwise influence are assessed with Pareto-smoothed importance-sampling leave-one-out cross-validation (PSIS-LOO).




Figure 3: PSIS-LOO stability. Histogram of Pareto- k values. Most $k < 0.5$ indicates stable importance weights and negligible high-influence points under the fitted hierarchy.

Residual diagnostics check for missed structure across covariates (e.g., $\ln g_{\text{bar}}$, stellar mass, surface density, redshift, environment) and for heteroscedasticity.




Figure 4: Residual checks. Standardized residuals versus $\ln g_{\text{bar}}$ and ancillary covariates, with binwise RMSE. No statistically significant trends remain after fitting, consistent with a single transition scale.

Generalization is evaluated with leave-one-galaxy-out (LOGO) validation.




Figure 5: Predictive validation. In-sample versus LOGO held-out RMSE by $\ln g_{\text{bar}}$ bin, and empirical coverage of binwise predictive intervals. Agreement indicates good generalization across galaxies.

Model comparisons among alternative response forms (kept under the same hierarchy) are summarized by $\Delta\text{ELPD}_{\text{LOO}}$ with standard errors and cross-checked against marginal-likelihood estimates.




Figure 6: Evidence summary. Relative support for the fiducial response (defined in S1) versus alternatives. Bars show $\Delta\text{ELPD}_{\text{LOO}}$ with ± 1 SE; dots indicate $\Delta \log Z$ where stable.

10.5 Supernova spill-over null tests

To exclude cross-contamination (“spill-over”) between distance-ladder cosmology and galaxy-scale dynamics, SN Ia Hubble-diagram residuals are fit with three contamination templates added to the standard cosmology:

1. Constant offset (homogeneous leakage)
2. Linear redshift trend (slow evolution)
3. Cubic spline with knots across the survey range (flexible, smooth structure)

Each template carries weakly informative priors and is compared via LOO and marginal likelihoods. Across all forms, the contamination amplitude is consistent with zero and shows no predictive improvement over the baseline—indicating no detectable spill-over from disk-galaxy kinematic phenomena into SN distances.

10.6 Priors & Hyperpriors

Table 1: Prior specifications

Quantity	Prior (suggested)	Notes
g_0	$\text{log-normal}(\ln g_0^*, \sigma_{g_0})$	Center g_0^* near calibration; choose σ_{g_0} weakly informative.
ν	$\text{normal}(\nu^*, \sigma_\nu)$	Center near the fiducial (e.g., ~ 0.59); avoid binding.
κ	$\text{normal}(1, \sigma_\kappa)$	σ_κ = external calibration budget (e.g., 0.02–0.05).
ϵ_i	$\text{normal}(0, \sigma_{\text{eps}})$	Galaxy residual; σ_{eps} has a half-normal hyperprior.
σ_{eps}	$\text{half-normal}(\tau_{\text{eps}})$	τ_{eps} set weakly informative.
Obs. scale	half-normal or half- t	Matches measurement model.
t dof (ν_t)	fixed or shifted-exponential	e.g., ν_t in 4–7; or $(\nu_t - 2) \sim \text{Exp}(\lambda)$.

Prior-predictive checks verify these priors do not induce spurious low- g curvature.

10.7 Convergence & Sampling

Fitting uses NUTS (e.g., 4 chains, X warm-up + Y post-warm-up draws per chain). All key parameters satisfy $\hat{R} < 1.01$ with bulk/tail effective sample sizes well above 400. Posterior predictive checks show calibrated coverage across $\ln g_{\text{bar}}$ bins.

10.8 Sensitivity & Ablations

- **Inclination / M/L drift:** broadened priors; transition parameters remain stable.
- **High-influence points:** refit after removing top 1–2% by Pareto- k ; changes well within errors.
- **Binning / transform:** alternate $\ln g$ binning and spline residual fits; no trend emerges.

- **Alternative response forms:** smoothed change-point, softened exponents; inferior ΔELPD but consistent transition.

10.9 Data & Code Availability

- **Rotation curves:** SPARC sample (175 galaxies; see main text).
- **SN Ia residuals:** standard compilation cited in the main text.
- **Code and figure scripts (S1–S6)** are archived at: [insert repo or DOI].
- **Environment:** [language/version], [probabilistic programming library/version].
Random seed(s): [insert].

11 Acknowledgements

We thank the maintainers of the SPARC database for making high-quality mass models publicly available, and the supernova community for open Hubble-diagram compilations that enable independent null testing. We are grateful to developers and contributors of PyMC and ArviZ for probabilistic programming and diagnostics that made the hierarchical analysis tractable. Constructive feedback from colleagues on calibration stability and predictive validation helped clarify the presentation.

This work used open-source software including NumPy, SciPy, pandas, matplotlib, PyMC, and ArviZ. Computations were performed on personal workstations; no proprietary data were used. Any remaining errors are our own.

12 References

12.1 Galaxy Dynamics & SPARC (Data/Background)

- Lelli, F., McGaugh, S. S., & Schombert, J. M. (2016). SPARC: Mass models for 175 disk galaxies with Spitzer photometry. *The Astronomical Journal*, 152(6), 157.
- McGaugh, S. S., Lelli, F., & Schombert, J. M. (2016). Radial Acceleration Relation in Rotationally Supported Galaxies. *Physical Review Letters*, 117, 201101.

- Milgrom, M. (1983). A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis. *The Astrophysical Journal*, 270, 365–370.
- Famaey, B., & McGaugh, S. (2012). Modified Newtonian Dynamics (MOND): Observational phenomenology and relativistic extensions. *Living Reviews in Relativity*, 15, 10.

12.2 Supernova Distance Data (Residuals / Compilations)

- Scolnic, D., et al. (2022). The Pantheon+ Analysis: The full data set and cosmological constraints. *The Astrophysical Journal*, 938, 113.
- Brout, D., et al. (2022). The Pantheon+ Analysis: SuperCal, photometric calibration, and systematics. *The Astrophysical Journal*, 938, 110.
- Betoule, M., et al. (2014). Improved cosmological constraints from a joint analysis of the SDSS-II and SNLS supernova samples. *Astronomy & Astrophysics*, 568, A22.

12.3 Bayesian Hierarchies, Robust Likelihoods, and Samplers

- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D., Vehtari, A., & Rubin, D. (2013). *Bayesian Data Analysis* (3rd ed.). CRC Press.
- Hoffman, M. D., & Gelman, A. (2014). The No-U-Turn Sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15, 1593–1623.
- Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, 27, 1413–1432.
- Vehtari, A., Simpson, D., Gelman, A., Yao, Y., & Gabry, J. (2019). Pareto smoothed importance sampling. arXiv:1507.02646 (updated version).
- Lange, K., Little, R., & Taylor, J. (1989). Robust statistical modeling using the t distribution. *Journal of the American Statistical Association*, 84(408), 881–896.

12.4 Model Comparison & Evidence Estimation

- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430), 773–795.
- Gronau, Q. F., Singmann, H., & Wagenmakers, E.-J. (2017). Bridgesampling for marginal likelihoods. *Behavior Research Methods*, 49, 884–897.
- Gelman, A., Hwang, J., & Vehtari, A. (2014). Understanding predictive information criteria for Bayesian models. *Statistics and Computing*, 24, 997–1016.

12.5 Rotation Curve Systematics / Mass Modeling (useful context)

- de Blok, W. J. G., & McGaugh, S. S. (1997). The dark and visible matter content of low surface brightness disc galaxies. *Monthly Notices of the Royal Astronomical Society*, 290, 533–552.
- Swaters, R. A., Madore, B. F., van den Bosch, F. C., & Balcells, M. (2003). The central mass distribution in dwarf and low surface brightness galaxies. *The Astrophysical Journal*, 583, 732–751.
- Lelli, F., McGaugh, S. S., Schombert, J. M., & Pawlowski, M. S. (2016). One law to rule them all: The radial acceleration relation of galaxies. *The Astrophysical Journal Letters*, 827, L19.

12.6 Cross-Checks & Residual Diagnostics (optional but referee-friendly)

- Hogg, D. W., Bovy, J., & Lang, D. (2010). Data analysis recipes: Fitting a model to data. arXiv:1008.4686.
- Gelman, A., & Hill, J. (2007). *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge University Press.